CHAPTER

12

TROUBLESHOOTING AND PROBLEM DETERMINATION

After reading this chapter and completing the exercises, you will be able to:

- Utilize sound troubleshooting logic to determine and solve problems
- Document problems and solutions
- Check for common causes of server failure
- Utilize network, connectivity, NOS, and hardware diagnostic tools
- Troubleshoot from a remote location
- Recognize and solve boot, virus, and hardware problems
- Locate help from vendors and peers

Even top-quality, optimally configured servers eventually develop problems, perhaps even serious enough to debilitate the server. In such events, it is critical that you implement sound troubleshooting logic to determine the exact problem and develop an appropriate solution. Sometimes the administrator becomes aware of server problems through user response to server performance or availability, and the wise administrator will lend some degree of credence even to nontechnical user comments. The administrator should also remain alert to direct and obvious causes of server problems or failure, and this chapter will remind you of some of the areas you should monitor.

Solving server and network problems goes much faster if you use the best tool for a given situation. Otherwise, you have to guess at the possible cause of a problem. Though there are dozens of server tools available that assist the administrator in troubleshooting, you will probably focus on a few favorites as well as the troubleshooting utilities that are included with the operating system. Sometimes you need to troubleshoot a server that is physically out of reach. In this case, some form of remote administration can be very useful in diagnosing and correcting server problems.

Most troubleshooting tools are available from within a functioning operating system. However, if you cannot boot the system, the possibility of recovering a failed server greatly diminishes. Recovering an unbootable system is critical in maintaining server health. For the most serious problems, many administrators restore a known good installation of the operating system using imaging software. Finally, when you cannot diagnose or recover a failed server, you should know where to turn to obtain the help you need, including help from the vendor or peers.

Troubleshooting and Problem Determination

Most of this book has already addressed one aspect of troubleshooting: prevention of problems by properly configuring server hardware and software. However, even with the best of servers, something will eventually go wrong for any number of reasons, including hardware failure, bad drivers, power problems, human error, and software conflicts, among other things. Sometimes there is no way to predict what a combination of hardware and software interactions will produce, and that's when trouble can arise.

Troubleshooting is one of the most frustrating parts of an administrator's job, but when successfully resolved, it is also one of the most rewarding. If you are used to troubleshooting home or workstation PCs, you'll find that the troubleshooting climate in the server room is much different than you are accustomed to. With a PC, troubleshooting success or failure affects only one person, and the pressure to successfully correct a problem is not as heavy. When a server fails, you might have hundreds of corporate users or online customers waiting on you. A single unavailable service or application can cost an organization thousands of dollars (or more) per minute. Hopefully, there is a level of redundancy that will continue to provide at least limited service while a server is down, but the seriousness of correcting server problems remains. During these times of intense pressure, it is critical to approach the troubleshooting puzzle with a logical, ordered perspective as shown in the remainder of this section.

Stay Calm

When catastrophe strikes and the pressure is on to fix a server, your pulse races, the phone rings, the pager sounds, and in the panic, it is easy to randomly stab at possible solutions before you've even confirmed the exact cause of the problem. However, it remains important to remain calm (as much as possible) and to follow a logical, step-by-step approach. The steps and guidelines offered here are not exhaustive, and you can insert steps or principles of your own according to your particular methods of troubleshooting. The important thing is that the administrator *has* a troubleshooting logic and methodology, and calmly pursues it.

Investigate the Problem

Investigation is usually the most significant step in troubleshooting, and it can also be the most frustrating. However, all problems give you at least a starting point. For example, if

the server won't turn on, at least initially you would suspect a power problem. Unless the cause of the problem is immediately apparent, the investigation process involves several possible stages, and your primary resource (besides the obvious symptoms of the server problem) begins with the log records. Investigation includes at least checking and keeping accurate log records, checking server messages, and asking analytical questions.

Document the Process

Documentation is a critical part of server administration. Otherwise, you are likely to unnecessarily repeat the same configuration errors or troubleshoot using the same failed methods as before. Good documentation starts with keeping good log records prior to the occurrence of a problem. When configuring server equipment and software, record exactly what you do and the success or failure of each step. Then, the log records have useful meaning when it comes time to troubleshoot.

The server also has records of its own, in the system logs and RAID logs, for example. There is a limit to how large these logs can get. Many RAID cards, for example, have a limited amount of memory to store the data. Be sure that you print out existing logs before clearing them to make room for new events. Server logs are stored on the hard disk but can become quite large. You can usually print these out too, or archive them to tape backup or offline media so that they don't occupy too much space.

With proper documentation, you are ready to use log records to quickly and efficiently troubleshoot the server.

Check Log Records

In other chapters, we discussed creating logs of any changes to the server or network electronically or on paper. I prefer electronic logs—that way you can access the log from anywhere on the network. The method you choose is not as important as the fact that you have a history of events that can affect the functionality of the server. Unfortunately, I know many administrators who do not keep logs because they are confident in their ability to remember their actions. However, this does not account for times when the administrator is unavailable, or in medium-sized or large networks where there are multiple administrators who might also need to know the history of a server. Documentation accounts for hardware assets and provides a progressive history that might reveal a series of actions that leads to a problem. Document every action performed on the server that could affect its functionality, including events such as:

- *Adding new peripherals.* Though external peripherals, especially a keyboard, mouse, and printer, might seem inconsequential, they can have a significant effect. For example, if a PS/2 mouse becomes unusable and an administrator replaces it temporarily with an old serial port mouse, a different set of resources will be used, and an IRQ assigned to the serial port will now be utilized, possibly conflicting with another device that previously accessed those resources. Also, these types of peripherals might include software that enables special features such as a

printer's double-sided printing capability. Any time you add software, you add hundreds or thousands of lines of code that could potentially interfere with other NOS functions. Peripherals such as new devices added to a SCSI chain become even more significant because of proper termination issues and ID assignments.

- *Installing software*. Although software packages interact with other packages and the NOS much better now than in the past, software bugs and interactions will always be potentially problematic. That's another reason why it is so important to test and pilot software deployments. I recently installed a file server with a CD burner and ATAPI stand-alone tape backup. Everything was fine until I installed the CD burner and tape backup software, which produced a wicked blue screen of death (BSOD) in Windows 2000. When I moved the devices to another server and installed the same software, there was no problem. These types of situations are somewhat out of your control, and you might not know that a problem could occur until testing it for yourself. Having documented that the CD burner and tape backup software produce a BSOD on that particular server will prevent me from making that mistake again.
- *Installing updates or upgrades*. One of the primary purposes of an update is to improve software and hardware compatibility. Nevertheless, some updates could cause more problems than they solve due to unforeseen incompatibilities. (In defense of programmers, it is nearly impossible for them to account for every possible software interaction that could present itself on a server.) In the example above regarding incompatible tape and CD burner software, I hope that the NOS or application vendors write a bug fix for the problem. If I apply such a patch, I'll enter it into the log records so that other administrators will know it is safe to install both applications.

Documentation of upgrades might also be important for proper license tracking.

- *Installing hardware and drivers*. In a Plug and Play NOS such as Windows 2000, allocating resources to various devices is much more flexible than in the past. Nevertheless, hardware devices will sometimes still conflict with one another. For example, the COM1 serial port typically uses IRQ 4, and many UPS systems connect to COM1. If you add a device that requires IRQ 4 and a power outage occurs, the UPS system might not be able to communicate with the UPS software through COM1. An equally common problem is hardware device drivers, which can cause any number of undesirable interactions such as incompatibility with the NOS, applications, or other devices. Note that while new drivers sometimes cause problems, updated drivers can also resolve problems.

Recall from Chapter 4 that there is a difference between UPS capacity (the volt-amps that the UPS supplies) and the UPS runtime (the amount of time the UPS can supply the volt-amps). People often confuse these two items. In many circles, you will still hear people mistakenly refer to extending the runtime as "increasing UPS capacity." Understanding what people really mean is a matter of understanding the context of the discussion.

Windows 2000 administrators will help to ensure device and driver compatibility by installing only drivers that are "signed" by Microsoft. Vendors send their devices drivers to Microsoft, where they are tested. If they are deemed stable and compatible, then a digital signature is included with the drivers. If the device is not signed by Microsoft, then Windows 2000 issues a warning like the one that appears in Figure 12-1, and you can deny the installation.

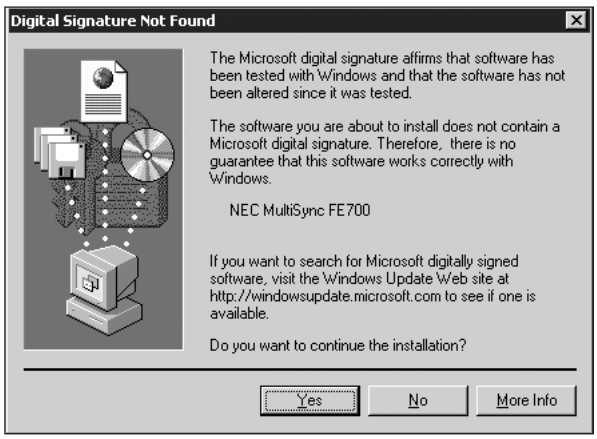

Figure 12-1 Windows 2000 notifies you of unsigned drivers

- *Interacting directly with other servers*. Many types of servers directly interact with other servers. A problem with one server can affect the functionality of any of the other servers. For example, some servers synchronize some type of information with other servers: mail, database, DNS, and WINS to name a few. If an authoritative DNS server were to have a connectivity problem, then you know that none of the other DNS servers to which it replicates will receive DNS updates until the connectivity problem is resolved, meaning that a problem with a single server can have widespread effects on other servers.
- *Stating server purpose*. Without documenting server purpose, various administrators can accidentally change a server's purpose beyond its original capabilities. Documenting a server's purpose can help to ensure that over time, the server is not re-purposed beyond its capabilities unless a compelling and deliberate reason

dictates otherwise. For example, you might want a mail server with more resources than it really needs to be used exclusively for mail so that it is as responsive as possible to mail functionality and, more importantly, future growth. If another administrator comes along and sees your shiny mail server with plenty of power to spare, he or she might be tempted to also use it to run another service or application. As the company grows with more new employees using email, you might be shocked to find that your more-than-capable mail server can't keep up because of its additional roles. Specially marking the server as dedicated to mail purposes might help to avoid this type of situation.

- *Identifying people performing work on the server.* If you run into a problem with the server and the documentation doesn't seem to help, the person who last performed the work might have additional insight into the cause of the problem. Similarly, it is important to identify persons to contact should a problem occur. For example, you are the administrator installing a new application on the server. Everything seems to work just fine, but you should document your name as the administrator who performed the installation. Better yet, enter the contact information of the vendor representative and/or technical support person.
- *Stating the purpose of the work on the server.* This allows other administrators to better understand the overall context of the server and the reasons for work that is performed on it.
- *Stating when the work started and finished.* This can help to develop a plan in performing similar tasks on other servers in the network, and can help pinpoint problems that might be attributed to the actions performed on the server.
- *Labeling cables.* Although it's not documentation in a log record, it is still prudent to label cables on both ends for easy identification. If you realize that the network cable to Server1 has a break in it, it's easy to find on the server end, but without labels on the other end, you might have to make random guesses as to which exact cable in the patch panel belongs to the server connection. Labeling is not always possible or practical, particularly in very large installations. Trying to locate the correct unmarked cable requires a Fox and Hound tool (discussed later in this chapter).

Another reason to keep records is for billing purposes. For example, if you visit a site as a consultant, I recommend that before you leave the site, you write down all actions performed and have someone initial or preferably sign the log to verify it.

Check for Server Messages

Fortunately, all major NOSs include at least rudimentary (though often cryptic) server messages to indicate the successful start or stop of services, various functions, server or software errors, system conditions, and more. For example, in Figure 12-2, this Windows .NET server returned an error that it was unable to find a domain controller to service logon requests. (In this case, the domain controller was down for maintenance.)

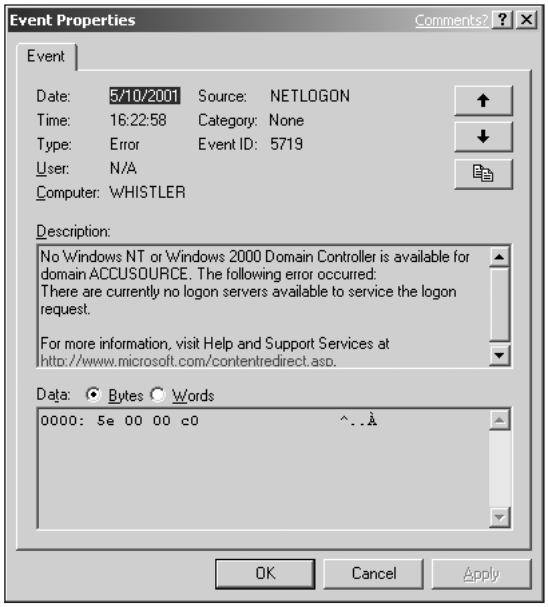

Figure 12-2 Windows 2000/.NET provides troubleshooting tips

Often, a Windows 2000 or Windows .NET event also provides tips on how to remedy the problem you encountered. For example, in Figure 12-3, you can see the beginning of one of four detailed recovery steps.

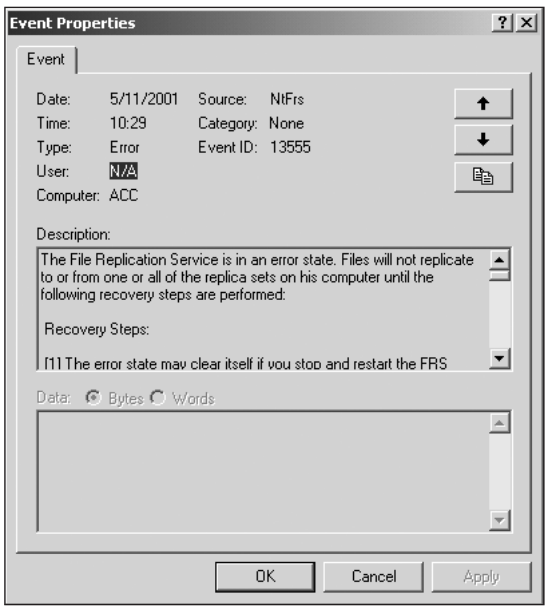

Figure 12-3 Event logs display a brief description of a problem

While writing this chapter, Microsoft revealed that the name of its new server product is Windows .NET, which is why we introduce it here. With Windows .NET, you can click a web page link in the Event Properties dialog box that, with your permission, sends system information to Microsoft that can be used to improve the Windows .NET product (the link is visible in Figure 12-2). Clicking the link takes you to a web page where you can perform research about the specific problem you encountered.

Often, event logs and error messages are cryptic and difficult to interpret. When stumped about what a message means, you can attempt to interpret its meaning by:

- *Referring to events that precede it*. Previous messages might indicate other services or functions that affect the message you are studying. For example, I recently saw a message on a Windows .NET server that indicated it was unable to retrieve a "backup list" after several attempts and that the backup browser was stopping. (The browser role in Windows networks enables you to view computers in My Network Places.) I read the preceding message and saw that the computer from which the server was attempting to retrieve the backup list was unavailable. Then I knew that the real problem was with another server and that if I could get the other server running, this problem would resolve itself.
- *Accessing the vendor's web site*. The NOS and some applications report error events in the NOS error messaging facility. Many error events are numbered—so if you visit the vendor's support web pages and perform a search for the event number, you might find a white paper or some other solution.

Although all error messages should be investigated, some may be innocuous, and as long as they do not affect performance, reliability, or availability, you can ignore them. For example, I regularly see the warning on a Windows 2000 server that appears in Figure 12-4. This occurred on several servers that I installed, many of them even before installing any applications or making configuration changes. I performed a search on "Event ID 3019" on the Microsoft Knowledge Base support site and found an article that explained the event. It's harmless. When you map a drive to local resources (as was the case with some user profile settings I had configured), Windows 2000 may be unable to determine a physical connection speed because it uses a software loopback adapter to locate resources on the same machine instead of using a true physical NIC; this is what generated the warning message. The article states: "This warning message is informational only and can be safely ignored."

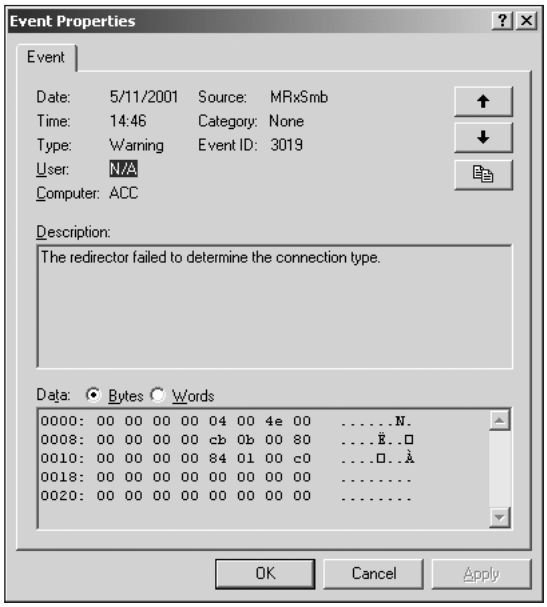

Figure 12-4 Some events are harmless and can be ignored

Aelita software offers several excellent administrative utilities, one of which is Aelita Event Management. This utility collects events generated from across the network into an SQL database, which can then be analyzed so that the administrator can see trends in Microsoft operating systems and applications as well as Novell NDS. For example, every time a user sends a print job to a Windows NT 4.0 or Windows 2000 printer, the print job is recorded in the event log. Using this software, you could collect print job information from several departmental print servers and analyze trends to see which departments have the fastest-growing print needs.

Exactly What Happens?

The important thing here is to track the root of the problem. Perhaps the problem initially presents itself through a few user calls saying that they are having trouble with email. Now, you have to ask a series of questions that narrows down the scope of the problem. Otherwise, you are likely to chase problems that don't exist. In determining exactly what happens in the email example, you might ask things such as:

- Are users having trouble sending, receiving, or both? If the user can only send, then SMTP over port 25 is probably working OK. If they can't receive, then perhaps another administrator accidentally blocked POP port 110 over the router, or the server's mail application has wrongly configured ports.

- What are the settings in the user's mail program? Though unlikely that each caller has identically configured the wrong mail settings, perhaps a deployment application that automatically configures the settings is incorrectly configuring each client. Or perhaps a new administrator spelled the name of the mail server incorrectly in the user's POP3 settings.
- Is the problem really with the server? Problems with users' email are a common occurrence and will often keep you pretty busy. However, user complaints about email are not always server related, and that's why you should ask specific questions about exactly what is happening. Most user issues revolve around incorrectly configured client email settings or network connectivity to their computer.

Do not neglect user input into emerging server problems. Though users are not typically server savvy, their descriptions can help you to narrow down potential server problems, either as they occur or prior to an emerging problem. For example, if a user states that Internet access is too slow, do not dismiss the complaint because you have a high-speed T-1 Internet connection. Perhaps there is a problem with the Internet proxy software, hardware, or both. If the proxy software usually caches Internet content to a particular hard disk that is currently failing, that would explain the slowdown in Internet performance.

How Does It Happen?

Much of the time, a problem occurs as a result of a sequence of events. When troubleshooting user clients, technicians ask the user what the user did to alter the system. (The answer is almost always "Nothing, I didn't change a thing.") With a server, there is usually nobody to ask if the event occurs during operations. That's when the log record comes in handy—so you can see if a previous change contributed to the problem. Also, look at logs generated by the application or NOS. They are often cryptic, but you might be able to distill enough useful information to determine a basic cause.

When Does It Happen?

Narrowing down the problem to a specific time or sequence of events helps you to determine what might be causing the problem. There are often scheduled events on the network that might contribute to the problem, or perhaps it's the result of peak traffic. Continuing the example of users with email problems, you might also ask questions such as:

- Does the problem happen when the users dial up to the network? If so, perhaps the problem is actually related to a RAS server and not email.
- Does it happen during a certain time of the day? Older backup programs cannot back up open mail storage unless it is first closed, which makes it unavailable to users.

- Does the problem predictably repeat itself, or is it intermittent? Problems that repeat predictably are much easier to troubleshoot because you can run performance monitoring applications or view logs to pinpoint when the problem occurs. You might also notice that predictable problems occur at the same time as another event. Intermittent problems are often baffling and difficult to diagnose because the problem might happen when you are not prepared to analyze its cause.

Intermittent hardware problems are often the result of ESD-damaged components, power problems such as power spikes or failing power supplies, or loose connections.

You must be sure that the amount of time you spend investigating the cause of a server problem is reasonable. Some organizations have a written administrator's guide that specifies the amount of time in which the administrator must resolve an issue before consulting other avenues, perhaps calling a vendor or referring the problem to another administrator. Do not allow your professional pride to prevent you from asking for help if you need it. Help might resolve the problem sooner and you can add to your own knowledge by learning from others.

Check the Obvious

It happens to every administrator eventually: You spend hours diagnosing the possible cause of a failed server using a full battery of tests, diagnostic software, technical support calls, and perhaps even a call to dial-a-psychic. Finally, you discover the only reason the server won't function correctly is because of an obvious problem that would have taken only a few seconds to remedy.

For example, a few years ago I had a friend (whom I'll call Keith) who worked as a Microsoft Exchange mail administrator for an organization of a few hundred users. The Exchange servers were well-tuned; however, users would periodically call complaining that they could not send or receive email. By the time Keith began to diagnose the problem, it seemed to fix itself and the user could again send and receive. This baffled Keith—he seemed to have done everything right in diagnosing the server but could not find any problems. Finally, he figured out what was causing the problem.

A few weeks earlier, the organization had hired a summer intern to work in the server room. Apparently, he had disconnected the mail server's RJ-45 jack from the patch panel and, in performing his duties, accidentally snagged the clip on the jack and it broke off. Nevertheless, he inserted it back into the patch panel, and apparently there was enough contact to allow normal communications most of the time. The server room was running too hot, and it became necessary to bring in an oscillating fan to temporarily assist the cooling. It so happened that when the fan would oscillate to the patch panel, it

would disturb the network cable enough to disrupt connectivity at times, and this accounted for the intermittent problems.

Although this scenario is not something you are likely to encounter, it does emphasize that you cannot forget troubleshooting measures that might at first seem too simple, such as *always check the physical connections*.

Physical Connections

It's easy to assume that because a troubleshooting problem is severe, the solution must be equally severe. There is no logical reason for this; it's just the way people sometimes react to problems in life and in technology. However, you can often save yourself hours (even days) of troubleshooting time if you check the obvious network cabling and physical connections first when troubleshooting hardware or connectivity problems.

Network Cabling

As the network grows, it becomes even more important to remember the importance of a network diagram, which is only as useful as it is accurate. For example, a growing network starts with five segments (three of which have connected nodes and two of which only connect the hubs) and four hubs. Adding another hub and segment to the network results in breaking the 5-4-3 rule discussed in Chapter 7, and you are likely to have very poor connectivity with high collisions. A network diagram helps you to keep track of network growth and avoid overextending its limitations.

With coax and twisted-pair cable, you might also run into a **bend radius** limitation, which impairs signal transmission when the cable is bent at too tight an angle. Typically, the cable should not bend more than four times its diameter to avoid signal loss. You are most likely to see exceeded bend radius where there is too much slack cable and it is all bunched up, or where the path of the cable run takes odd twists and turns.

Also remember not to exceed the maximum cable length for each respective cable type as discussed in Chapter 7.

Connections

Similar to Keith's loose RJ-45 connector mentioned at the beginning of this section, all kinds of connectors can come loose unless they include some kind of retention such as thumbscrews, latches, or locks. RJ-45 connectors are particularly problematic due to a clip that easily snaps off or becomes overextended and weakened. Don't forget to add a boot to the ends of patch cables to protect the clip. Also check component connections. USB connectors are usually tight enough to stay by themselves, but moving other cables and equipment around the connector might accidentally wiggle a USB cord as well, causing it to loosen. The same applies to FireWire connections and serial, parallel, and video connectors—though with the latter items you can use thumbscrews to securely attach the cables.

Inside the case, I have often encountered loose connections, usually with hard disk ribbon cables. Because of the 18-inch limitation of IDE cables, sometimes the cables really stretch to go from host adapter to the actual drive in full-size cases. I recently found an IDE CD-ROM that stopped working. Because of continuous tension on the cable, the connector had come partially loose. Changing the position of the CD-ROM drive in the drive cage so that it was closer to the IDE host adapter solved that problem.

Similarly, I also recently found a file server that only provided intermittent access to the ATAPI tape drive. The connections were tight and correctly oriented, but when the computer would POST, the drive only showed up about half the time. Since no operating system or applications were involved in the POST phase of the boot, the problem was more likely physical. Finally, I realized that the drive seemed rather far away from the IDE connector on the motherboard, yet the cable reached it just fine. After pulling out the cable and measuring it, I immediately realized the problem: Somebody got the bright idea to use a special 24-inch IDE cable because the drive bay was too far away for an 18-inch IDE cable. This meant that the signal strength was compromised and was the reason for this problem. Again, I rearranged components, used an 18-inch cable, and corrected the problem.

The width of SCSI and IDE ribbon cables often impedes good airflow inside the case. You can try to "flatten" the cables as best you can, but inevitably you'll still have airflow problems. I recently found SCSI and IDE cables from *www.coolerguys.com* that are round instead of ribbon (see Figure 12-5). Replacing ribbon cable with round cable can significantly improve airflow. Also, it is easier to daisy-chain round cable. Sometimes you have to twist ribbon cable because the number 1 pin is on the right on one device and on the left on the next device; this makes the distance between connectors in ribbon cable shorter and sometimes strains the connectors.

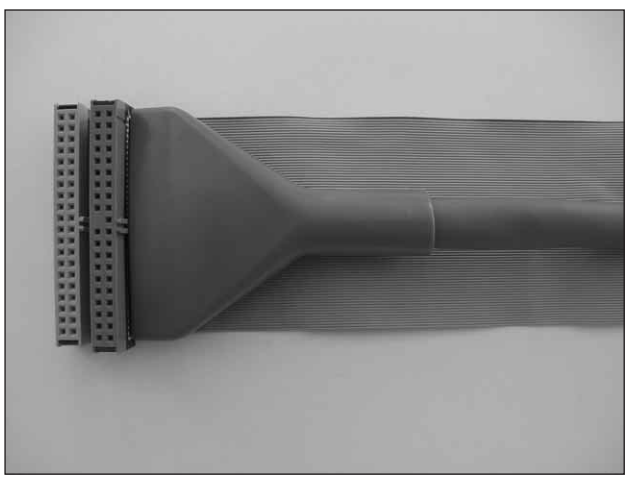

Figure 12-5 Round cables are narrower than ribbon cables, and allow much better airflow inside the case

Come to Your Senses

Your senses can often alert you to problems, hopefully in a proactive way to avoid trouble instead of as a reaction to disaster. You won't have to use all your senses to troubleshoot every situation. (After all, who "tastes" a server?) However, it is important not to ignore sight, sound, smell, and touch when troubleshooting.

People tend to depend on sight as the primary sense, and hopefully you can observe when something is wrong. With sight, just remember not to take anything for granted and to stay alert. Assuming all server hardware is correctly installed and configured, you would particularly want to observe any changes. For example, after sliding a server to the front of the rack, do any of the cables in the back come loose? Even with cable arms, cables sometimes loosen. Consider it a good practice to press in each connection when you reinsert the server.

Listen for unusual noises too, which usually come from the hard disk or a fan. When the hard disk starts to make unusual noises, failure is imminent and you should utilize redundancy measures such as RAID-1 or tape backup and immediately replace the drive. The noise is the result of mechanical failure, and it is normally not worthwhile to attempt to open the case and repair the drive. IDE drives are very cheap to replace, and expensive SCSI disks normally have a warranty program promising overnight delivery of replacements.

When you hear noise in the power supply or case cooling fans, it is usually a result of failing bearings (which under normal conditions provide an extremely smooth, low-friction action). Failure to replace the fan could cause the system to overheat (plus, the noise is really annoying). Replace case cooling fans from your stock of spare parts. With the power supply, replace the entire power supply, as its design is not intended for serviceability. Failure to replace a PSU can result in an overheated power supply, causing it to fail. Though the risk is slight, a fire hazard is also present.

Fans accumulate dust more quickly than any other single component of the system. Unusual noises might be the result of excessive dust impeding normal operation; see if blowing out the dust eliminates the sound.

Properly designed server rooms have very clean air and are probably one of the healthiest places for people to work. There will probably never be cigarette smoke in the room, and even food and drink are usually against company policy. If you smell anything unusual in the server room, it could be an administrator working long shifts without a shower. The other most likely smell will probably be the plastic-like smell of overheated components or the unmistakable smell of electrical smoke. The human sense of smell is likely to be able to detect these odors prior to a smoke alarm—so go into bloodhound mode and try to locate the source, and then immediately power down the problem equipment. You may be able to catch an electrical fire in the early stages and prevent a halon dump and the resulting mess.

Finally, the sense of touch will help you to detect problems such as an unusual amount of heat for potentially overheating components such as the power supply.

DIAGNOSTIC TOOLS

When troubleshooting server problems, several potential points of failure might be to blame. In order to properly diagnose the problem, the administrator must be aware of various diagnostic tools and their capabilities. The health of the network, while not the main focus of this book, is directly related to the effectiveness of the server, so the basic network concepts discussed in Chapter 7 are very important to remember. Each respective server NOS has accompanying network tools, and third-party vendors also offer a rich assortment from which to choose. Although it's not practical to detail each tool, this chapter addresses several examples of such tools for diagnosing problems with the physical network and network traffic.

Likewise, the administrator might suspect a problem with a physical device but find it difficult to confirm. For example, the administrator suspects a defective memory chip because of recurring blue screens. How can the administrator be certain that the problem is with physical memory and not with a bad driver or application that crashes into other processes in memory? For these types of problems, you need a thorough analysis utility in addition to following the troubleshooting principles in this chapter.

Network Cable Diagnostic Devices

When the network stops, so does your business. The common misconception is that most network issues are difficult to properly diagnose. The truth is that with proper training, finding a network problem is usually a matter of using proper diagnostic equipment discussed here and the troubleshooting methods discussed in the first section of this chapter.

As discussed in Chapter 6, a multimeter combines the functions of a voltmeter (which measures the potential difference, or voltage, between two points) and an ohmmeter (which measures resistance) as well as a few additional functions.

One common application of using a multimeter is to diagnose and troubleshoot network errors on a typical thinnet or 10Base2 network. Recall from Chapter 7 that this cabling requires 50 ohm resistors placed on either end of the bus to properly reflect the signals down the wire. Should a resistor or a T-connector fail, or an individual segment fail, the entire cable segment will fail as well. By using a multimeter, you can effectively troubleshoot by tracking each segment's resistance.

Also, you can use a **time domain reflectometer (TDR)**, a tool that not only detects cable breaks but also provides the approximate distance to the break by measuring the time it takes for a signal to return. A TDR does not tell you exactly what the problem is, but it will tell you where the problem is located on the physical cable.

If it interests you, there is a formula to calculate the distance to the problem section of a cable. A signal is sent down the cable, and the TDR measures the time it takes for the signal to return, converting time to distance, which is then divided by the speed of light and multiplied by the proper velocity of propagation (VOP). (VOP is a measure of the speed of light multiplied by mitigating factors that affect this speed, such as the physical cable media. Twisted-pair cable has a VOP of about .65.) Finally, the result is divided by two. The resulting number is the distance to the problem location on the cable.

In addition to checking for cable breaks, also check cable lengths. If you exceed cable lengths only slightly, connectivity problems might appear periodically, but not necessarily on a consistent basis. Similarly, inconsistency in network signaling could occur if you are within maximum cable lengths, but other factors such as radio signals or EMI compromise signal integrity.

Another device in your network diagnosis arsenal is a **Fox and Hound**, or **tone generator and locator**, used to identify cable. Imagine pulling several hundred cables through walls, raised floors, and conduit from the patch panel. Identifying each individual cable ahead of time would be fruitless, as the labeling might not survive the trip. By sending a tone, or a signal, down a cable segment, the Fox and Hound can easily identify each cable, and then you can appropriately label the cable. A Fox and Hound is actually a pair of network tools—a tone generator that applies a tone signal to a wire pair or single conductor (the "fox") and an inductive amplifier locator probe (the "hound") on the other end. At a cross-connect point such as a patch panel, or even at the remote end, you can use the amplifier probe to identify the conductor within a bundle to which the tone has been applied.

The Fox and Hound can also be utilized to troubleshoot physical problems with the cables. For example, some tone generators will also allow you to test resistance levels and provide audible tones to indicate the line condition.

Connectivity Utilities

To test for network connectivity on TCP/IP-based networks, the first tool most seasoned administrators will turn to is the Ping utility, which verifies N etwork layer connectivity. (The Network "layer" is a reference to the third layer of the OSI network model. If you are not familiar with this model, refer to *Network+ Guide to Networks* by Tamara Dean, from Course Technology, ISBN 0-7600-1145-1). Recall from Chapter 7 that the Ping utility verifies remote host accessibility by sending small packets of data to which an accessible host responds. This functionality takes place by sending an "echo request" to a remote host such as another computer or router. If the destination host does not respond, the interface will display some form of an error message, such as "destination host unreachable."

To determine if a host is properly configured on a TCP/IP network, execute the following sequence of steps:

1. Ping the IP address (or host name) of a network device that lies outside of your local router's interface (commonly referred to as your "default gateway").

If the destination host replies, the system is properly configured. If the destination host is unreachable, it is time to locate the issue: Is it a local configuration issue, or is it on an intermediate system?

2. Ping the IP address of the local router's interface to your subnet (default gateway). If this IP address replies, the issue lies outside of the local subnet. Perhaps the destination host is offline, or there is a network issue somewhere in between. If the default gateway is unreachable, determine if the problem lies with the router or in the local system.
3. Ping a local host on the local subnet. If this host replies, the problem most likely resides with the router. Perhaps the router's interface into the LAN is down due to an administrative error, temporary maintenance, an error condition, or because the router itself is in the process of rebooting. Also check the IP address of the default gateway in the system's configuration, as it may be wrong. If the local host does not reply, you probably have an issue with the local system.
4. Ping the local host's IP address. A reply indicates you may have misconfigured your subnet mask. If you do not get a reply, ping the loopback address of 127.0.0.1. If you get a reply, you probably mistyped your IP address when configuring TCP/IP properties. If you do not get a reply, it is time to check the TCP/IP stack, which is often remedied with a simple reinstallation of the protocol. Finally, remember to check physical connectors and drivers:
   - Is the network cable connected to your NIC?
   - Is the network cable plugged into the hub or switch?
   - Are the proper drivers loaded and running?

The loopback address is commonly known as 127.0.0.1, but the entire 127.x.x.x range is set aside as loopback addresses. For example, you could also enter 127.255.255.254 to Ping the loopback address.

You can also use several command switches to extend the functionality of the Ping command, as shown in Table 12-1.

Table 12-1 Useful Ping Command Switches

Switch	Purpose	Use it when . . .
-t	Ping the specified host until interrupted.	The standard four packets do not return data for an acceptable length of time. For example, you could run it continuously to identify when a router starts "flapping," which is becoming intermittently available/unavailable. (Linux Pings indefinitely by default.)
-a	Resolve addresses to host names.	You want to verify the presence of DNS reverse lookup records.

Table 12-1 Useful Ping Command Switches (continued)

Switch	Purpose	Use it when . . .
-n *count*	*count* represents the number of echo requests to send.	You want more Ping packets than in a typical Ping, but do not want to Ping indefinitely.
-l *size*	*size* represents the size of the buffer.	You want to test how network equipment handles packets of various sizes.

In past years, malicious users were able to disable servers by sending what is known as the "ping of death," which is a ping –l packet with a very large buffer. NOSs such as Windows NT 4.0 could not handle the large packet and would become disabled. Microsoft and most other NOS vendors have resolved this issue with patches, though many other similar denial-of-service attacks continue to be a serious issue.

With NetWare 5.1, you can also use the TPing command, which is very similar to Ping.

The Ping command is not case sensitive in most operating systems, but remember that UNIX/Linux is case sensitive, so you will need to type "ping" in all lowercase to use the utility. Also, Ping is not a separately loaded program—it is an integral part of the TCP/IP protocol stack. In other words, if you have TCP/IP loaded and functioning correctly, you also have Ping. There is no such thing as a "ping service" or "ping daemon." If you can otherwise access the server (logging in, for example) but cannot successfully ping the server, this indicates that there is probably a problem with the TCP/IP stack on the server, and you will need to reload TCP/IP by rebooting. If the TCP/IP stack is corrupt, you will need to reinstall it.

Another commonly used network diagnostic utility is TRACERT (or TRACEROUTE). Depending upon the operating system or hardware platform, this utility traces the route your packets take to reach the destination host on a TCP/IP network (as discussed in Chapter 7). You may use this utility to determine if there is a network issue between your host and the destination host. The resulting output will show the path of the packet from one hop to another. This may diagnose a slow router, a dead router, or a misconfigured access list.

To speed up TRACERT, you can use the "-d" switch to stop resolving IP addresses to their associated host names. You may also wish to set a maximum trace length by minimizing the maximum number of hops to the target host by using the "-h *maximum_hops*" switch.

Operating System Utilities

Many operating system utilities have already been addressed in Chapter 11, "Performance Monitoring," because performance monitoring is a big part of troubleshooting the server. For example, Chapter 11 mentioned that a very high frequency of unaccounted-for interrupts usually indicates that a hardware device has failed. Performance monitoring tools detect such events. However, there are also utilities that you can use specifically for troubleshooting, many of which are listed in the sections that follow.

OS/2 Utilities

Table 12-2 shows OS/2 utilities that you can use to monitor system health and performance.

Table 12-2 OS/2 Utilities

Utility	Purpose
PSTAT	Displays specific status information about individual processes. For example, it can show threads and DLLs involved in a process as well as memory being shared by other threads.
RMVIEW	Locates and manages hardware resources (IRQ, DMA, I/O, and so forth).
RESERVE.SYS	Allocates resources for devices that are not Resource Manager aware. Resource Manager is an OS/2 function that automatically allocates hardware resources similar to, but not as comprehensive as, Plug and Play.
SystemView	Remotely manages servers and automatically upgrades software.
SMBTool	Captures and displays network traces.

If OS/2 loads a driver that causes problems, you can easily identify which specific driver is to blame by pressing Alt+F2 during the boot when you see a "??? OS/2" message in the upper left of the screen. OS/2 will then display the name of each device driver as it loads, but be aware that the problem driver is often the second-to-last driver (not the very last) displayed on the screen.

System hangs under any operating system are difficult to handle, because many times you can only guess if the entire operating system has hung or if it is just the operating system interface that has hung (while the operating system continues to run beneath it). With OS/2, you can use the system clock to determine the true state in most cases. Open the system clock and look at the second hand. If it is running, then only the desktop has hung, and you should consider waiting a while longer to see if a task finishes. If the second hand has stopped, the entire OS has stopped.

Novell NetWare Utilities

Table 12-3 lists NetWare utilities that you can use to monitor system health and performance.

Table 12-3 Novell NetWare Utilities

Utility	Purpose
VREPAIR	Repairs a volume.
CONLOG	Captures console messages to a text file for later viewing.
NWCONFIG	Modifies the server configuration, performs management operations, and installs additional products.
WAN Traffic Manager	Manages how and when WAN traffic is sent.
TPCON	Monitors TCP/IP activity.
NCMCON	Controls and monitors hot-plug PCI devices.
DSREPAIR	Repairs NDS database problems.

Besides these utilities, you can continue to use NetWare Administrator (NWADMIN or NWADMN32) or ConsoleOne to perform general administration.

Linux Utilities

Table 12-4 shows graphical (KDE or GNOME) Linux utilities that you can use to monitor system health and performance.

Table 12-4 Linux Utilities

Utility	Purpose
Tripwire	Detects changed files and directories.
KDE Control Center	Provides detailed system information about the system's applications, devices, and GUI interface.
Sysctlconfig	Configures specific settings for networking, file systems, virtual memory, and the kernel.
KDE Task Manager	Similar to Windows NT/2000 Task Manager, you can view each running process and adjust its priority level ("re-nice") or kill it.
tksysv and SysV Init Editor	A system process editor that allows you to add and delete services as well as adjust priority (see Figure 12-6)

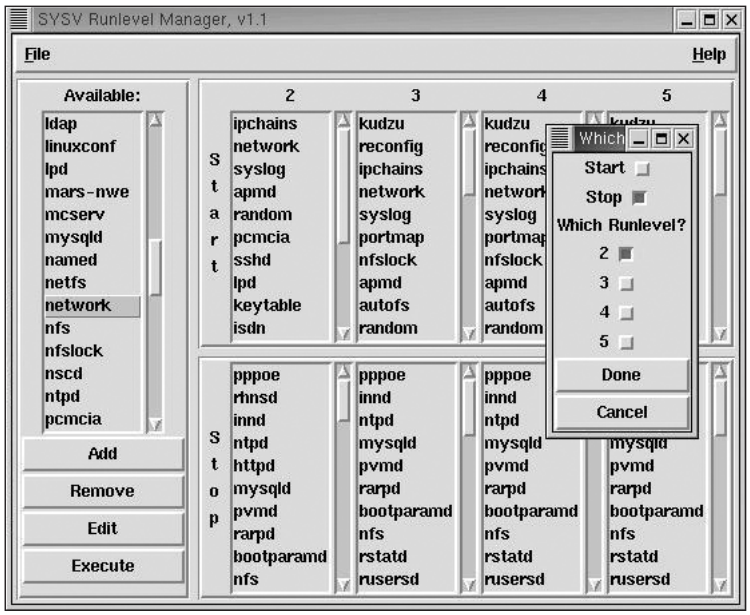

Figure 12-6 The Linux tksysv utility allows you to manipulate various processes

Windows NT 4.0 Utilities

Table 12-5 shows Windows NT 4.0 utilities that you can use to monitor system health and performance.

12

Table 12-5 Windows NT 4.0 Utilities

Utility	Purpose
Disk Administrator	Manages disks and partitions; creates software RAID configurations.
Task Manager	Shows running applications and processes, and allows you to adjust their priority or terminate. Displays real-time processor and memory utilization statistics.
System Properties	Adjusts how applications use memory, copies/deletes user profiles, configures virtual memory, and manages the boot menu.
Windows NT 4.0 Resource Kit	A collection of dozens of various utilities designed to make administration, troubleshooting, and performance monitoring more effective.
Windows NT Diagnostics	Accesses data regarding the system bus, BIOS, CPU(s), display adapter, memory usage, paging file usage, running services, system resources, and more.
Dr. Watson	Debugger for win32 applications. Might not make much sense to you unless you're a developer. Software vendors might ask for a copy of a Dr. Watson output to analyze their programs.
Network Monitor	A basic network sniffer useful for packet analysis.
The /SOS switch	Add this switch to the end of the Boot.ini file to view each driver as it loads, similar to the OS/2 Alt + F2 boot method.

Windows 2000 Utilities

Table 12-6 shows Windows 2000 utilities that you can use to monitor system health and performance.

Table 12-6 Windows 2000 Utilities

Utility	Purpose
Computer Management	Broad management capability that includes event viewer, system information, device management, hard disk management, various services, and other functions.
Task Manager	Shows running applications and processes, and allows you to adjust their priority or terminate. Displays real-time processor and memory utilization statistics.
Windows 2000 Resource Kit	A collection of dozens of various utilities designed to make administration, troubleshooting, and performance monitoring more effective.
Dr. Watson	Debugger for win32 applications. Might not make much sense to you unless you're a developer. Software vendors might ask for a copy of a Dr. Watson output to analyze their programs.
Network Monitor	A basic network sniffer useful for packet analysis.
The /SOS switch	Add this switch to the end of the Boot.ini file to view each driver as it loads, similar to the OS/2 Alt + F2 boot method.
Active Directory Domains and Trusts	Establishes trust relationships between domains.
Active Directory Sites and Services	Manages Active Directory replication between sites.
Active Directory Users and Computers	Creates and manages Active Directory user, group, and computer accounts.
System Information	Comprehensive information about system hardware, components, and software.

System and Hardware Diagnostic Utilities

Because much of troubleshooting is actually a matter of ensuring proper configuration in the first place (and correcting misconfigurations), using the right administration tool is important. Administrative tools for the various operating systems were discussed briefly in Chapter 8.

Besides the administrative utilities included in the NOS, performance monitoring utilities, as discussed in Chapter 11, are also critical because troubleshooting often involves improving the performance of a server. Sometimes a diagnostic utility is included with an installation package for a device as seen in the screenshot for the Intel NIC (see Figure 12-7).

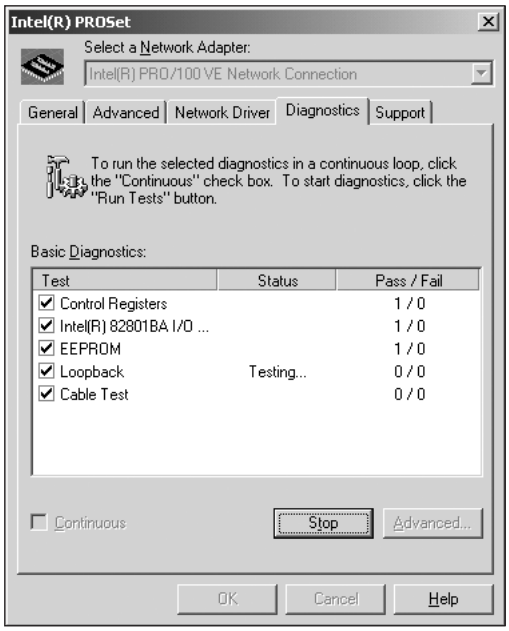

Figure 12-7 Running diagnostics on an Intel NIC

Most server vendors include utilities that perform low-level diagnosis of the server hardware. You can usually place these utilities in a directory and run them from there, or copy the utility to a floppy disk. Floppies are useful for systems that cannot boot properly, or for file systems such as Windows NT NTFS that do not allow direct disk access from a conventional MS-DOS boot disk.

Third-party manufacturers have also made utilities that diagnose hardware. There are several such products, many of which are designed for individual workstation use. One such product that I have found useful on servers is American Megatrends' AMIDiag, which performs extensive diagnostics for workstations or servers. Most utilities of this type can run only from MS-DOS, but AMIDiag offers a convenient and intuitive Windows interface as well. The latest version includes loopback plugs for serial and parallel port testing, and reports detailed information about components, including the motherboard, chipset, memory, and processor. For server components you suspect might be faulty, you can run extensive diagnostics as shown in Figure 12-8.

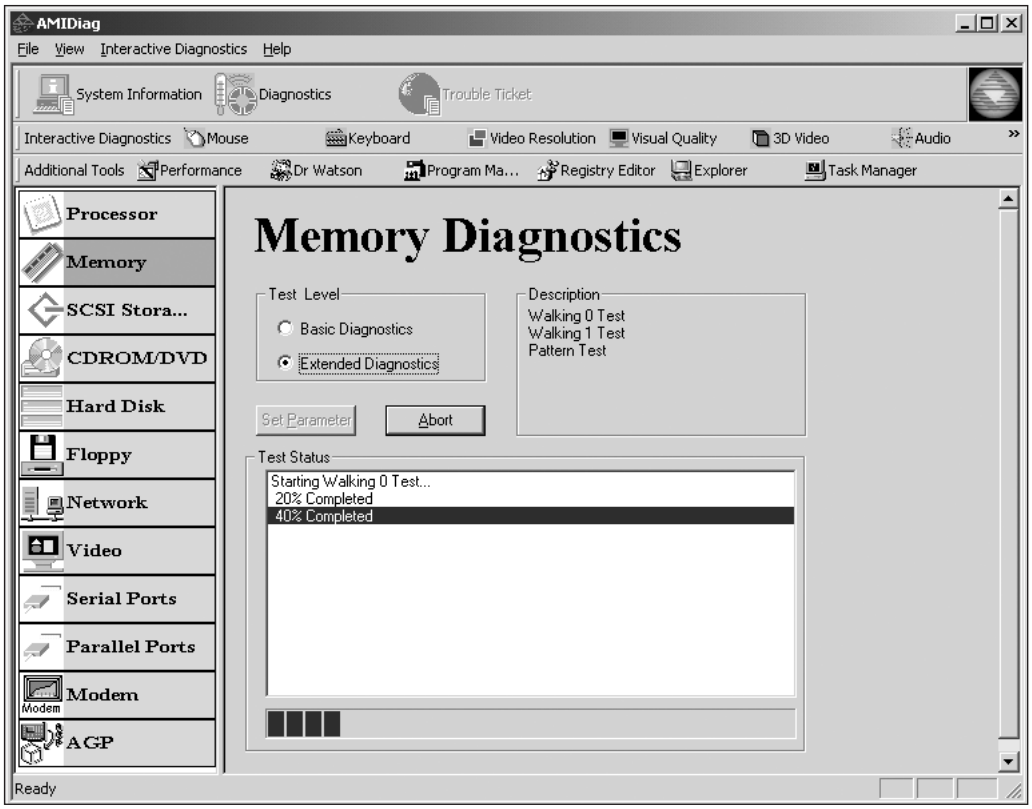

Figure 12-8 AMIDiag performing extensive tests on system memory

In some situations, you might also consider creating a small FAT partition on the hard disk that can be used for several purposes, especially when you need to perform administration on a system without loading the NOS. Typically, loading the NOS causes a problem because the file systems of most NOSs are inaccessible from MS-DOS. You might store common MS-DOS utilities in the partition that are useful for disk management, moving or editing plain text files, and so forth. For example, I would start with:

- ATTRIB—Without ATTRIB to reset attributes on some files, you might not be able to view or change certain files.
- FDISK—For repairing the master boot record using the /MBR switch or adding additional partitions. In the worst case, you will use FDISK to delete partitions.
- FORMAT—To format partitions after using FDISK.

- CHKDSK or SCANDISK—To check the health of FAT partitions.
- MSCDEX—To load CD-ROM capability, also include the CD-ROM drivers.

The partition can also be useful for storing diagnostic software. For example, if the system won't boot, you can't access NOS diagnostics, and the CD-ROM is inaccessible as well, meaning that third-party diagnostics are also unavailable. By storing an MS-DOS diagnostic utility in advance on a separate partition of the hard disk, you save yourself the headache of trying to work around the inaccessible NOS file system.

Working Remotely

Unless you enjoy working in the tight confines of a server or telecom closet at all hours of the day or night, you will quickly learn to embrace working remotely. Wake-on LAN network adapters, remote management features on the NOS, and third-party remote management software allow the administrator to work from nearly anywhere to administer and troubleshoot the server.

Wake-on LAN

Although most servers run 24/7, organizations that run mostly during business hours might choose to conserve power by putting servers and workstations into a low-power state during off hours, with only certain servers (such as web or RAS servers) running constantly. If the administrator wants to perform administration or troubleshoot during off hours, he or she must typically wait until server utilization is low, which translates into a long day in the server room for the administrator. However, with various remote features such as wake-on LAN, you can administer servers from nearly anywhere (such as the comfort of your own home.)

Wake-on LAN (WOL) is a technology that allows the administrator to remotely wake a computer from its low-power state. WOL allows administrators to "wake up" the server and perform tasks during times of reduced activity. You can use WOL on both Token Ring and Ethernet networks.

WOL works when a **magic packet** (or **wake-up packet**) consisting of 16 copies of the MAC address is sent to the host system from a server system, which has a remote network management application installed. When the WOL NIC receives the magic packet, the server turns on. To enable WOL, you must have the following:

- A WOL network interface card. This card is always awake, listening for the magic packet. In order to do this, the adapter card draws continuous, low power from the motherboard.
- A WOL motherboard BIOS
- WOL remote management software

WOL functionality has several advantages, even for client workstations. For example, every workstation in your organization can be started up just before your employees arrive, saving the time and revenue lost while they boot, and avoiding unnecessary energy costs associated with employees leaving systems on all night.

Remote Administration Tools

There are many utilities and tools available that allow you to remotely administer the systems in your network, whether they are across the room or in another state. One of the most widely used is Symantec's pcAnywhere (*www.symantec.com*), which allows you to remotely connect to your server to copy files, run remote applications, or perform any other actions that you would normally do if seated locally at the server.

Another free method of remotely administering servers is Windows NT 4.0 or Windows 2000 Terminal Services. In Windows 2000, use Add/Remove Programs to add Terminal Services. A wizard asks you to specify the mode in which you want the server to run:

- *Application server mode* allows users to run the Windows 2000 desktop and applications from any Windows workstation, even Windows 3.x. The benefit is that you can use older hardware incapable of running Windows 2000 or more demanding applications on its own. Using an older version of Windows, clients can derive all the benefits of a Windows 2000 desktop. Application server mode requires purchase of proper client licenses.
- *Remote administration mode* allows up to two administrators to connect as if local to the server at no additional licensing expense (see Figure 12-9). Of course, this is the mode we are interested in for purposes of remotely administering a server.

Either mode allows an administrator to "shadow" an ongoing session. You can see what applications the connected terminal services user is using and, if you like, operate their keyboard and mouse. For user sessions, this can be an excellent support tool. As a server administration tool, Terminal Services allows you to perform complete administration over a LAN, the Internet, or a dial-up connection. For example, if you are out of town and your pager receives an alert that a service has unexpectedly shut down, you could dial up from your hotel room, establish a terminal session, and restart the service. If the situation is drastic, you could even remotely reboot the server.

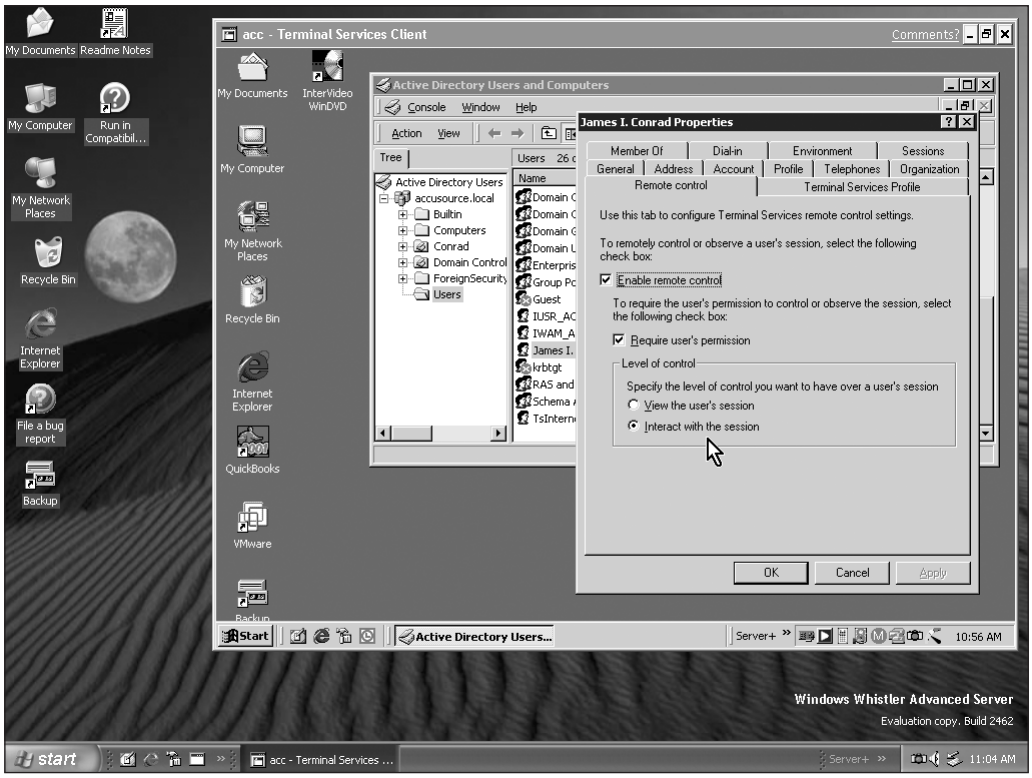

Figure 12-9 Remotely connecting to a Windows 2000 Advanced Server computer from Windows .NET Server

OS/2 has a similar remote administration tool OS known as SystemView Remote Control.

Another way that administrators can connect to and troubleshoot remote desktops and servers with a simple right-click is through the Computer Management console, as follows:

1. From any Windows 2000 desktop, right-click the My Computer icon.
2. Choose Manage. The Computer Management console appears, and by default connects to the local computer.
3. Right-click the Computer Management node at the top of the left pane, and select *Connect to another computer*.
4. Select the computer of your choice from the interface shown in Figure 12-10, and click OK.

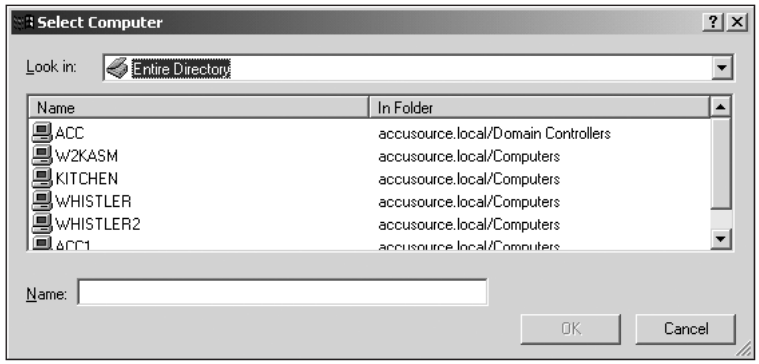

Figure 12-10 Select the computer that you want to remotely administer

5. The Computer Management console displays management nodes for the remote computer (see Figure 12-11) and functions for the most part as if you were locally seated at the server. Some features are disabled; for example, you can view Device Manager hardware settings, but you cannot change them unless seated locally or running over a Terminal Services session.

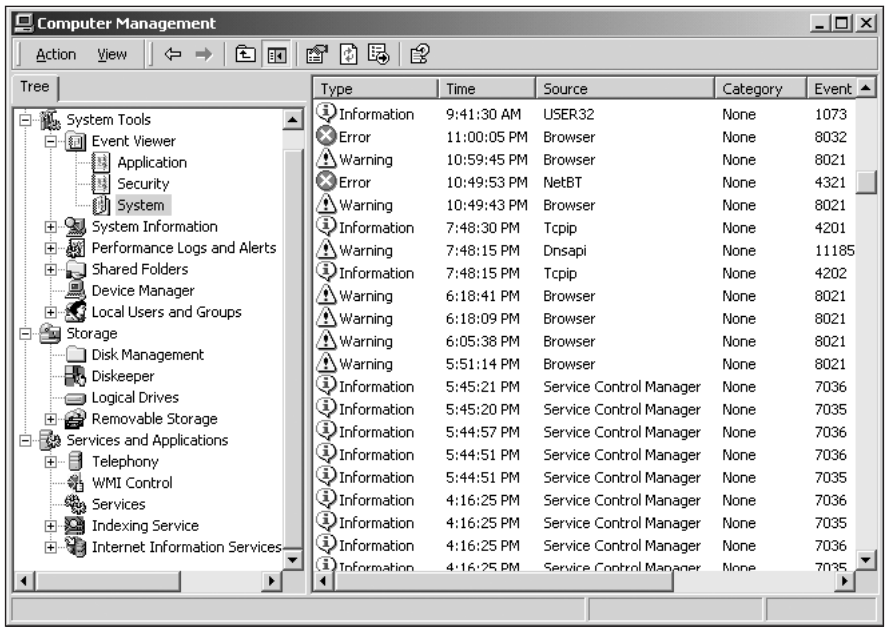

Figure 12-11 Remotely administering a Windows 2000 computer

The above steps apply to Windows 2000 and Windows .NET, but most NOSs offer a similar remote management facility. For example, Windows NT 4.0 has a similar capability, and NetWare has ConsoleOne and Remote Management Facility for general administration

and ZENworks for remote control capability. You can also use the RConsole or RConsoleJ. With OS/2, use the SystemView utility (also known as NetFinity).

In addition to software products, there are hardware solutions that provide a great deal of remote administration flexibility. These are usually PCI or ISA cards that occupy an available slot in the server and are connected via a network cable, internal or external modem, or both. One of the best products I have seen is MegaRAC (Remote Access Companion) from American Megatrends. Because of an onboard battery and WOL support, you can administer the server remotely even in the event of a power failure. Some other remote administration tools require the server to be up and running, but this is of little use if you need to remotely change BIOS settings or observe startup error messages, text sequences, and so forth. With MegaRAC, you can still perform all of these actions from a remote station. Some of the features of MegaRAC are as follows:

- Captures the screen when the server crashes so that you can see error messages or what might have caused the problem. This is especially handy for error messages that might appear briefly before an unplanned reboot such as some instances of a BSOD. In addition, it logs the probable cause of the crash.
- Monitors system health including temperatures, fan RPMs, voltages, and alerts for chassis intrusion.
- Allows you to view the graphics display of a remote server to see current activity.
- Allows you to reboot a server and watch the boot process (including POST).
- Provides a special boot to a separate partition that loads diagnostic utilities. You can also insert a floppy in your local computer, and use it to boot the remote server! Very handy for booting an unbootable server, or running utilities or diagnostics that would otherwise require you to visit that server.
- Provides dead server management. When the server fails, MegaRAC can issue an alert to the supervisor and put itself in a special server mode to receive commands from the supervisor, such as performing a reboot or running special diagnostics. This is probably the single most powerful benefit of the product.

Booting an Unbootable Server

At some point the inevitable will happen—the server will refuse to boot. Numerous recovery options are available in each NOS, but here we primarily address basic boot functionality to provide you with a starting point for troubleshooting.

Typically, the first option is to uninstall any new hardware devices, device drivers, or software. Windows NT/2000/.NET allows you to boot to the "last known good" state, which bypasses changes made by devices/drivers/software in the last session. If the system still refuses to boot, try using a boot disk. Most operating systems come with one or more

bootable floppies or CDs to use during setup. It is always wise to guard these diskettes closely and make backup copies. Instructions for returning to a basic bootable operating system for each respective NOS appear below.

To create boot floppies for a Red Hat Linux system from MS-DOS or Windows:

1. Change directories to the CD disk drive where the installation files are located.
2. Change to the dosutils directory.
3. Execute the *rawrite* command, which then prompts you for the file name of a disk image. Consult your Linux documentation for a suitable boot image. In Red Hat Linux, for example, use the \images\boot.img file on the Red Hat CD-ROM.

Recall from Chapter 8 that there are several boot images, each with various purposes, such as a conventional Linux boot or booting to a network.

4. Specify the destination letter for the floppy disk (usually "A"). The contents of the image file copies to the floppy.

To create boot floppies for a Red Hat Linux system from Linux:

1. Mount the CD with the installation files.
2. Change directory to the desired image.
3. Execute the following command: *# dd if=boot.img of=/dev/fd0 bs=1440k.*

Create boot floppies for an OS/2 system in any of the following ways:

- Use the Create Utility Diskette utility in the System Setup folder.
- From the command prompt, use the *makedisk* command.
- From the installation files OS2\INSTALL directory, run the Bootdisk.exe program. The disks allow basic access to the hard disk and include simple functions such as a text editor, FDISK, format, and so forth.

To create the boot floppies for a Windows NT-based system:

1. Have three labeled floppy disks ready.
2. Enter *rdisk* from the command prompt, and follow the onscreen directions that automatically appear, swapping disks as prompted.
3. Please note that the disks will be inserted in reverse order; that is, disk 3 will be inserted first, and disk 1 last.

The floppies that the RDISK utility creates are not directly bootable, but contain information that will allow you to run setup in "repair" mode to restore critical boot and system information.

To create the boot floppies for a Windows 2000 system:

1. Have four floppy disks ready.
2. On any Windows or MS-DOS computer, insert the Windows 2000 CD into the CD-ROM drive.
3. Execute the *makeboot a:* command from the *\bootdisk* directory of the CD, and follow the screen prompts to switch disks.

Booting to Windows NT or 2000 with the boot floppy disks does not provide full access to the operating system; it only allows you to use other recovery tools to repair the Registry or add/remove files. Prior to a server boot failure, however, you can create a special (unsupported) boot floppy to access the NOS desktop as follows:

1. Format a floppy disk from within the running Windows NT/2000 NOS. This is important, as it creates a boot sector that would not appear if formatted under MS-DOS or Windows 9.x.
2. Copy the following files to the floppy:
   - BOOT.INI
   - NTLDR
   - NTDETECT.COM
   - [NTBOOTDD.SYS] (if a SCSI adapter is present)
   - [BOOTSECT.DOS] (if the computer dual-boots)
3. Now, provided there are no other NOS problems, the system should boot as usual to the desktop.

Boot disks for many operating systems may be found at *www.bootdisk.com*.

Recall that NetWare actually uses a conventional DOS partition to boot. If the boot files are corrupt, you can boot with any bootable MS-DOS floppy and manually create or edit Autoexec.bat and Config.sys files, which service the basic MS-DOS boot. Server boot files include Autoexec.ncf, Startup.ncf, and Config.ncf. As a reference for the correct commands to include, you can copy known good files from another properly functioning NetWare server.

NetWare requires at least 10–20% free disk space on the SYS volume to boot reliably. If you suspect disk space is an issue, run "NDIR /VOL," which will show total volume space that is in use and is free. If space is an issue, run "PURGE /ALL" for each volume, which permanently deletes previously erased files (similar to emptying the Recycle Bin on Windows operating systems). If you still haven't opened enough space, you may want to run FILER or NWAdmin to manually remove files.

More Drastic Measures

Finally, if server problems are so severe that you cannot recover within a reasonable length of time, you might have to:

- Reinstall over an existing operating system. This is probably a better method than formatting the hard disk and starting over, because reinstallation usually replaces files that might have become corrupt, and applications usually remain intact. I recommend this method in situations where you are certain that there are missing or corrupt NOS files but cannot identify each one and manually replace it within a reasonable length of time. This method is a bit more of a gamble if you are less certain that NOS files are missing or corrupt. If the problem was a virus, for example, reinstallation will probably not prevent the virus from again damaging the NOS. Some NOSs such as Linux also offer special routines that are designed to run a setup program specifically to repair an installation.
- Use cloning software. Alluded to earlier in this book, you can use imaging software to store the entire contents of a partition into a single file that you can later reapply to the same or a different computer. Create an image of a known good NOS installation, so that if problems occur at a later time, you can reapply the known good NOS image and effectively roll back time. I prefer this method when possible, because it restores the server in minutes and you do not have to install applications or spend time configuring services and NOS settings.
- Format the hard disk and reinstall the NOS and all applications, and then configure all services. While this is not too disruptive for a workstation and might take between one and four hours, on a highly utilized server, this could take at least a day and be highly disruptive. If this is the method you use, server logs are again crucial: You would not want to duplicate a series of events that led to the server corruption in the first place. Also, by looking at the recorded purpose of the server, you avoid running services and applications that are unnecessary.

Troubleshooting Viruses

I'll assume we all know that viruses are programs that are sometimes harmless pranks, but more often are self-propagating programs that perform damaging and wide-ranging actions on server and client computers. The most effective way of troubleshooting viruses is to use a preventive, constantly running virus-detection utility that recognizes the presence of a virus and cleans infected files if possible. Unfortunately, if a virus gets loose on a system, there is little that virus utilities can do to undo the damage. A few virus symptoms include:

- Unbootable or intermittently bootable systems. Many viruses, especially early viruses, attack the boot sector.

- Display problems. Some make screens unreadable—although they may or may not damage actual data.
- File destruction. Obviously this is the most comprehensive damage, though some viruses might also slightly alter data and, therefore, are more difficult to notice.
- Unexplained poor performance (for example, slow program load times or disk access times)
- Unexpected low-memory situations
- Strange (even rude!) noises and/or graphics
- Unexplained reduction in available disk space
- Random drive letter reassignment
- Odd error messages
- A growing number of bad hard disk sectors
- Mail servers handling an extremely high number of emails. Some particularly destructive viruses send copies of themselves to all persons in the client mail program's address book.

If a system becomes infected by a virus, you might be tempted to delete and re-create all partitions. If you then boot from a floppy and reinstall the OS from a CD, you might unwittingly reintroduce the virus if the floppy is also infected. Be sure to scan the boot floppy for infection prior to using it.

All major virus detection/protection software should work fine for you; the key is to make sure that the virus definition files that identify signs of a virus are up to date. I prefer to use McAfee's web-based VirusScan Online, which checks for updates automatically on a daily basis, though all these antivirus programs are excellent:

- McAfee VirusScan (*www.mcafee.com*)
- Norton Antivirus (*www.symantec.com*)
- Innoculan (*www.ca.com*)

For more about viruses and their effects, refer to Jean Andrews' book, *Enhanced A+ Guide to Managing and Maintaining Your PC* (Course Technology, 2000).

Troubleshooting Specific FRUs

Recall from earlier chapters that an FRU is a Field Replaceable Unit. Depending on the context, an FRU can refer to the server as a whole or to the components themselves; for example, the power supply is a common FRU.

This section addresses some of the most common symptoms of FRU failure, while Chapters 3, 4, 5, and 6 addressed identification of specific parts and how to replace them.

Power Supply

Unless the system utilizes N+1 power supply redundancy, loss of a power supply means total loss of server function. Signs of a failing power supply include:

- Intermittent memory problems. Because clean, consistent power to the memory ensures its integrity, power problems can cause memory loss or corruption. If you regularly see errors that report a problem in the same memory address, the problem is more likely the actual memory.
- Systems that lock, hang, or reboot. While I was writing Chapter 8 of this book, one of my servers suddenly hung without any apparent reason. A few hours later, it hung again. I realized that juggling five internal hard disks, two CD-ROMs, and a tape drive all at once for the past couple of years was finally too much. And yes, I should have sized the power supply better. Anyway, systems with overtaxed or failing power supplies sometimes exhibit these problems, especially when they get busy.
- Damaged or intermittently failing motherboards. Power irregularities are bad for all electrical components; the motherboard is no exception. Continuing the lesson learned in the previous bullet, I'd like to announce that I am the proud owner of a new motherboard.
- Unusually hot case and power supply. Minutes and seconds without adequate cooling can damage the power supply, and failing fans are often the cause.

Use the backprobing techniques discussed in Chapter 6 to verify proper power supply voltages.

Memory

Memory problems can often be difficult to diagnose. Operating systems and applications are so complex that when a crash occurs, it might be difficult to say for certain whether it was due to a software bug or memory problems. However, the following might help you to diagnose memory problems:

- *Varying part numbers*. Some cheaper memory is really made from manufacturer spare parts. In fact, some memory modules are actually a collection of chips from different manufacturers. In this case, the markings on the chips will all be different. However, recall that it is normal for one chip to have different markings if it is the parity chip. Memory timing is so tight that little if any variation can be tolerated. For server use, do not use modules with varying part numbers.
- *No POST*. There may be no POST if memory is not properly seated, the memory is of poor quality, it uses mixed parts (as described above), or incompatible modules are installed (ECC and non-ECC in the same server, for example).

- *Memory errors in the same location each time*. Most servers include a reporting facility within the CMOS that will show the memory error. If not, and you have the opportunity to otherwise record the memory location, make a note of it. Later, if a memory error appears again in the same location, you know it is because of the memory module, not because of the NOS or applications. However, some device drivers might require the same location in memory, and a bad driver might be disguised as bad memory.
- *Memory-testing utilities*. I again recommend AMIDiag, though there are several utilities that can do this for you. The utility might run from within the NOS, but I suggest running it from a bare boot floppy with no memory management in effect (such as Himem.sys or loadhigh statements).
- *Physical symptoms*. Bad memory chips will often run comparatively hotter than the others. Also inspect the system board slot contacts. If a contact is bent or misaligned, you might have to replace the motherboard.
- *Suspect modules.* When replacing suspect modules, replace them one at a time and test them each time to confirm which modules are bad.
- *Handle with care.* Exercise special care in handling memory modules, which are particularly susceptible to ESD problems.

Hard Disks

Some hard disk problems are obvious—especially when they make that grinding noise that sounds like there's gravel in the drive. If the master boot record is damaged or corrupt (often due to virus infection), the system may not boot. Try booting from a floppy to verify that the system is boot capable, and to rule out other possible problems. Other hard disk problems include data loss or corruption, and unusually high or growing numbers of bad sectors as reported by disk analysis utilities.

Use the FDISK /MBR switch to reconstruct the master boot record for Windows, MS-DOS, or any other NOS that starts from a DOS-compatible boot partition. Similarly, Windows 2000 allows you to boot to a special recovery console, which allows you direct disk access with the right password. Then, you can use additional boot utilities such as FIXMBR and FIXBOOT.

If you experience hard disk problems, consider the following solutions:

- Perform a thorough disk analysis of the hard disk. All major NOSs have such a utility included, such as Windows CHKDSK command. Also consider third-party utilities, but remember that the more thorough the utility, the longer it will take to analyze the hard disk (perhaps hours). Most NOSs also have hot-fix capability to automatically save data away from detected bad sectors onto healthy disk space.

- Verify proper cabling. For IDE, make sure that the marked Pin1 on the cable is adjacent to the power connector. Also make sure that the power connector is inserted all the way. SCSI connections are a study of their own (refer to Chapter 5 for proper configuration information and the next section for specific troubleshooting issues). Verify proper termination, signaling (HVD, LVD, or SE), and device ID settings.
- Check the POST. ATAPI drives will display their presence during POST, and SCSI adapters such as Adaptec's will show the specific drives in the setup BIOS of the adapter.

SCSI and RAID Troubleshooting

As alluded to in Chapter 5, troubleshooting SCSI is a study of its own. This section addresses these issues specifically, and though some information might overlap Chapter 5, it warrants a solid review. Recall from earlier in this chapter that troubleshooting is mostly configuring devices properly in the first place, and it could never be more true of any device than it is of SCSI drives. Adding to the obvious importance of properly configuring SCSI as an administrator, note that SCSI is heavily emphasized on the Server+ exam!

The following subjects resolve the most common SCSI issues.

Termination, Termination, Termination!

If someone tells you that the new SCSI drive is not working, you can almost assume that it's a termination problem. Even experienced administrators occasionally forget to properly terminate the SCSI chain. Consider the following SCSI termination issues:

- Remember that both ends of a SCSI chain need to be terminated.
- If the drive does not support self-termination, then you must add a terminator to the end of the chain. (The cable might already have a built-in terminator.)
- Most newer drives such as Ultra160 drives do not self-terminate. Some older 50-pin drives offer self-termination.
- If there is no room for a terminator, use a pass-through terminator that both connects the drive to the chain and provides termination. Even better, use a cable that has a terminator already crimped to the end.
- Passive termination is generally considered the worst type of termination because its resistance (132 ohms) varies too greatly from the standard cable resistance of 105 to 108 ohms. Make sure you use active termination instead of passive termination. Even better, use forced perfect termination whenever possible.
- Recall that when adding 50-pin devices to a 68-pin bus, you must specially terminate for the additional pins that the 50-pin device does not use. Although there are 18 additional pins, nine are only for grounding and do not require special treatment. The remaining nine, known as the "high nine," are hot and require you to properly terminate.

Cabling

- Recall from Chapter 5 that 8-bit ("narrow") SCSI always uses 50-pin cables and connectors, and 16-bit ("wide") SCSI always uses 68-pin cables and connectors. Therefore, if you are combining various generations of SCSI technology on the same bus, make sure the cabling remains compatible. For example, Ultra2 SCSI devices (which are usually "wide" devices) use a 68-pin connector and so does Ultra3. In this case, you could mix the technologies.
- If the Ultra2 SCSI device is the more rare "narrow" version requiring a 50-pin connector, then you cannot use it on an Ultra3 68-pin bus without an adapter.
- You can use special adapters to place 50-pin devices on a 68-pin cable.
- Likewise, you can use special adapters to place 68-pin devices on a 50-pin cable.
- This is important: If you mix 50-pin narrow devices with 68-pin wide devices, the net effect is that the 68-pin devices are effectively reduced to 50-pin narrow performance. There are all kinds of converters, custom cables, and so forth that allow you to physically mix and match devices and buses in countless ways, but conversion usually comes at the cost of performance.
- If you have 50-pin and 68-pin devices on a 68-pin bus, place the 68-pin devices nearest the host adapter. Then add and terminate the 50-pin devices.

If you need to buy SCSI cables, terminators, or adapters, I strongly recommend *www.cselex.com*. Not only is there equipment for most every feasible SCSI combination, but the technical support persons are well-informed and responsive.

Device ID

Recall that each device has a method to set the device ID, usually through jumpers as is often the case with hard disks or with a dial.

- Especially on older SCSI, set the booting hard disk to ID 0 or else it might not be recognized as the boot drive.
- If you also have IDE drives in the system but want to boot from SCSI, then configure the system BIOS to boot from the SCSI host adapter as the first bootable device. Some systems may require you to connect the IDE drive to the secondary IDE interface instead of the primary IDE interface.
- Each SCSI device must have a unique ID number. Recall that in order from highest to lowest, the priority order is 0, 7, 6, 5, 4, 3, 2, 1, and then if the bus is also wide, 15, 14, 13, 12, 11, 10, 9, 8.
- The SCSI host adapter is usually set at 7; it is best to leave it at this setting.

RAID

Keep the following RAID tips in mind:

- A RAID log records the state of the RAID array. This log can be in the host adapter BIOS, the system log of the NOS, or separate software.
- With RAID-5, you must keep the drives in the same order in which they were originally configured. For example, if a remote office sends you a pre-configured RAID-5 external array case but takes out the drives and ships them separately, you must place the drives back into the original locations on the SCSI chain. If the drives are placed in a different order, the RAID controller will return a message that the array configuration has changed and that you should reconfigure it.
- Do not move drives configured with one RAID adapter to another RAID adapter. RAID host adapters actually write data to track zero or the last track on the drive, and this information is often largely proprietary. Moving the drives could cause the RAID array to not recognize the drives or worse, destroy data.
- As a rule, any time you "touch" equipment that contains data, back up first. For example, if you are asked to add an additional drive to the RAID array, back it up first! No exceptions.
- The host adapter BIOS automatically implements the hot spare when a RAID drive fails. The interim time between when the original drive fails and the hot spare is fully implemented is known as "degraded mode." For example, if a RAID-5 member fails, it takes time for the array to integrate the hot spare while it reconstructs data from the parity stripe of the other array members.
- If you want fast performance without redundancy, use RAID-0.
- If you want fast read performance and low overhead in terms of usable storage space in the array, use RAID-5. RAID-5 lags behind other RAID methods for write performance because of the parity calculation.
- If you want fast read and write performance at the expense of high disk overhead, use RAID-1.

Adapter Cards

Although it may seem obvious (if the device stops working, it's defective!), there are still a series of questions and solutions you should look into for NIC problems.

- Move the card to another computer. If the problem repeats itself in a known good system, then you have isolated the card as problematic.
- Closely associated with the physical hardware is the device driver. Upgrade the driver to verify that the driver is not corrupt and that it is the most compatible.

Usually, the most recent version will include bug fixes for problems of which the vendor is aware. (You updated a NIC driver in Hands-on Project 6-7.)

- Move the device to a different slot. Some devices perform better when in a different slot, and high-performance devices work nominally better when closer to the processor.
- Check to see if the device is fully seated into the slot.

Processor

The processor might fail, but this is somewhat rare and is usually brought on by overheating. Usually when a processor fails, the system won't even POST. Motherboard failure is also rare but not unheard of, and it is usually brought on by power problems or overheating. If you suspect a motherboard problem, use an extensive diagnostic utility to analyze the board.

TROUBLESHOOTING TIPS

Finally, here are some general troubleshooting tips based on my own experience.

- *It usually takes longer than you think.* When you know what is causing a problem, it is easy to quickly review the steps that it will take to resolve it in your head. However, it is easy to sometimes leave out steps that you don't realize until you're in the thick of a repair. Also, in resolving one problem, other parts of the system are often affected.
- *Start with the simplest solution.* I alluded to this earlier in the chapter, but it is so true. Save time and headaches by starting with the simple causes and solutions. Also, more complex troubleshooting steps tend to affect the rest of the system more significantly. For example, it is much less disruptive to uninstall the most recent application or hardware device than to reinstall the operating system.
- *Move hardware to see if the problem follows.* If you suspect a specific hardware device is problematic, move it from the original system to a known good system and see if the problem repeats itself. If so, you have confirmed the hardware problem. If not, you know that the problem is something else. It might also be possible that the hardware device interacts poorly with other hardware or software in the original system.
- *Always check for updates.* I recently installed a new software package on a Windows 2000 server, and upon reboot, the system returned a BSOD. About ten days later, the vendor came out with a patch for the exact problem I experienced.

- *Troubleshoot one issue at a time, and apply one solution at a time.* The tendency is to attempt too many solutions at once. If you think a problem has three possible solutions, do not perform all three at once. Execute each solution individually to see if it is the proper solution. That way, you can enter the exact solution in the server log.
- *Fix or change one thing at a time.* Similar to the above issue, it is not wise to try the "kill two birds with one stone" method of troubleshooting and maintenance. For example, if you install a new hardware device and driver, add new software, and apply an operating system patch all for the same session, how will you know which one caused a problem if the system does not properly function afterward?
- *Progressively document troubleshooting steps.* I like to keep a legal pad nearby, and write down the history of steps I take to troubleshoot an item. This record is unofficial, and once a successful solution appears, I'll enter it into the actual server log. Of course, any of the steps that change the system permanently must also be entered in the server log. The purpose of the history is if the troubleshooting session gets lengthy and tiresome, it is easy to forget whether or not a particular solution was already attempted. In addition, it is easier to backtrack a specific sequence of steps if they make the problem worse or add new problems. I also make a note each time the system successfully reboots. If the server does not reboot successfully, then I can look at the series of steps that led up to that point.

  For example, I recently installed a brand new server with all applications. I made a hasty mistake of installing too many apps at once, and upon reboot, the system returned an unrecoverable error. Because I had installed several applications and put off the requisite reboot for each one until I was done with all the applications, I had no idea which application caused the error, or if it was a combination of the applications. Fortunately, prior to the applications, I had used disk imaging software—so I was able to return to a known good state. Then, I performed each installation separately and rebooted after each application. In case you're wondering, the problematic software was a CD-burning software application, which did not seem to operate well with other applications.

Getting Help

Administrators have a tendency to be somewhat independent and determined to single-handedly solve server ills. If the administrator is solving a problem on a workstation or even a home PC, that approach might be OK. But when the network is at the mercy of a failed server, you can't take too long to resolve the issue. One of the best ways to quickly resolve these problems is to get outside help. You will also want to get help if determining the cause of the problem and its solution requires more knowledge than you currently have. It is not good timing to learn a new technology in the midst of a server disaster; therefore, you should consult a source with more information. Somebody has probably

already solved the problem you're dealing with; there is no need for you to "reinvent the wheel" on your own. Help is available from a number of sources such as:

- *Web support*. This is usually the first place to go, because most vendors that really want to support the customer will keep searchable white papers and other technical support documents. Using web support also avoids those annoying automated phone menus and repeating messages about your value as a customer. The web is also the main source for product updates, a grateful change from the days when administrators had to order and pay for updates to be mailed on a floppy disk.
- *FAQs*. A listing of frequently asked questions (FAQs) can appear on the application, NOS, or hardware installation disk. While you might not read the entire thing, it might be worthwhile to at least scan it to see if any of the issues it raises will apply to your situation. An FAQ list might also appear on the vendor web site.
- *Phone help*. Contact telephone technical support from the vendor of the product with which you are experiencing a problem. Although it seems that phone technical support persons can sometimes be much less knowledgeable than you are overall, they probably have enough specific knowledge about their product to at least offer some suggestions.
- *Newsgroups*. There are two main types of newsgroups. The first is vendor moderated, which keeps the comments civil and apropos. The second is public and unmoderated, which might be less reverent about the vendor but might also be more truthful. In both cases, you use the newsgroup reader to track a problem to see what the vendor suggests as well as what other administrators might be able to contribute. These can be excellent sources of peer help—chances are, somebody else has experienced a troubleshooting problem similar to yours and can offer a solution. The disadvantage of newsgroups is that they are often not instant solutions. You might post a message today and have to wait several days for an answer, if one returns at all. Similar to the newsgroup is the chat room in which others contribute their knowledge to solving your troubleshooting problem.
- *Email*. Many vendors offer email support. Typically, you go to the support section of the vendor's web site, and if your problem is not solved there, a link might appear for you to contact technical support via email. Sometimes you will receive a response the same day, and sometimes never. I recently received a response to a query I sent six weeks earlier.
- *Newsletters*. Various email newsletters are an excellent source of late-breaking product news and optimization and troubleshooting tips. Newsletters are distributed on a periodic basis—usually daily, weekly, or monthly. For example, Red Hat has several newsletter mailing lists available at *www.redhat.com/mailing-lists*, and an excellent Windows 2000 newsletter can be found at *www.trainability.com*.

Chapter Summary

- By properly configuring server hardware and software, you largely avoid server problems. However, sometimes there is no way to predict what a combination of hardware and software interactions will produce, and problems arise. That's when troubleshooting begins.
- It is easy to randomly stab at possible solutions before you've even confirmed the exact cause of the problem. However, it's important to remain calm (as much as possible) and logically approach the problem.
- Whether you keep paper or electronic logs, it is crucial that you keep accurate log records so that others can learn the history of the server, avoid past problems, and retain the server's original purpose (unless a compelling and deliberate reason dictates otherwise). The log also records server interaction with other servers, changes relating to applications, updates, and drivers, and the persons performing work on the server.
- All major NOSs include at least rudimentary (though often cryptic) server messages to indicate the successful start or stop of services, various functions, server or software errors, system conditions, and more. Go about interpreting cryptic error messages by referring to events that precede it or accessing the vendor's web site.
- In analyzing server problems, ask yourself "Exactly what happens?" "How does it happen?" and "When does it happen?"
- Server problems often occur as a result of obvious, easily overlooked problems such as physical connections (network cable, network connections) as well as signs that your sense of sight, smell, sound, or touch can detect.
- When troubleshooting server problems, several potential points of failure might be to blame. In order to properly diagnose the problem, the administrator must be aware of various diagnostic tools and their capabilities. Each respective server NOS has accompanying network tools, and third-party vendors also offer a rich assortment from which to choose.
- Important network diagnostic devices include a multimeter, Fox and Hound (tone generator and locator), and a time domain reflectometer (TDR). Besides the hardware tools, you still have the ever-handy TCP/IP utilities such as Ping and TRACERT. With Ping, you generally ping a device outside the local router's interface, ping the IP address of the local router on your subnet, ping another local host on your local subnet, ping the local IP address, and finally, ping the loopback address (127.x.x.x).
- Useful OS/2 utilities include PSTAT, RMVIEW, RESERVE.SYS, SystemView, and SMBTool.
- Useful NetWare utilities include VREPAIR, CONLOG, NWCONFIG, WAN Traffic Manager, TPCON, NCMCON, and DSREPAIR.

- Useful Linux utilities include Tripwire, KDE Control Center, Sysctlconfig, KDE Task Manager, tksysv, and SysV Init Editor.
- Useful Windows NT 4.0 utilities include Disk Administrator, System Properties, Windows NT Diagnostics, and the /SOS switch.
- Useful Windows 2000 utilities also include Computer Management, Active Directory Domains and Trusts, Active Directory Sites and Services, Active Directory Users and Computers, and System Information.
- Both Windows NT 4.0 and Windows 2000 include Task Manager, respective resource kits, Dr. Watson, and Network Monitor.
- The performance monitoring utilities discussed in Chapter 11 are also critical, because troubleshooting often involves improving an underperforming server. Some devices include diagnostic utilities. Most server vendors offer some type of diagnostic utilities that you can use to perform low-level diagnosis of the server. You can also use third-party diagnostics to perform thorough hardware analysis.
- Remote administration allows you to administer servers regardless of physical location. Several technologies allow this functionality, including Wake-on LAN network cards, third-party remote administration software, and remote administration hardware.
- When a system won't boot, start by uninstalling any new hardware devices, device drivers, or software. Windows NT, 2000, and .NET allow you to boot to the "last known good" state, which bypasses changes made by devices, drivers, or software in the last session. Also, try using NOS utilities to create a bootable floppy disk.
- If you cannot recover a server in a reasonable amount of time, you might have to take one or more of the following more drastic measures: Reinstall over an existing operating system, use cloning software to restore a known good instance of the system, format the hard disk and reinstall the NOS and all applications, and then configure all services.
- Viruses are a common cause of server failure, and you should use antivirus software to detect and clean infected files. There are many symptoms of virus infection, including rapid and comprehensive file deletion.
- This chapter includes tips for troubleshooting and replacing an FRU such as power supply, memory, hard disk, and adapter cards. It also includes the following troubleshooting principles: Troubleshooting usually takes longer than you think; start with the simplest solution; move the hardware to see if the problem follows; always check for updates; troubleshoot one issue at a time and apply one solution at a time; fix or change one thing at a time; and progressively document troubleshooting steps.
- Obtain help from other sources when it is taking too long to resolve the problem or you need more information about a technology. Helpful sources include web support, FAQs, phone help, newsgroups, email, and newsletters.

Key Terms

bend radius — A limitation with network cable that impairs signal transmission when the cable is bent at too tight an angle. Typically, the cable should not bend more than four times its diameter to avoid signal loss.

Fox and Hound (or **tone generator and locator)** — A pair of network tools. The tone generator applies a tone signal to a wire pair or single conductor and, using an inductive amplifier probe (locator) on the other end, will permit you to identify that conductor within a bundle at a cross-connect point such as a patch panel, or even at the remote end.

magic packet (or **wake-up packet)** — A packet consisting of 16 copies of the MAC address sent to the host system from a server system, which has a remote network management application installed. When the WOL NIC receives the magic packet, the server turns on.

time domain reflectometer (TDR) — A network troubleshooting device that measures the approximate distance to cable breaks.

wake-on LAN (WOL) — A technology that allows you to remotely wake a computer from its sleep mode. WOL works by sending a "magic packet" from a remote station to the WOL host. The "magic packet" contains 16 copies of the WOL host's MAC address.

Review Questions

1. Which of the following is not a valid reason to keep a server log record?
   a. to identify and contact persons who have previously worked on the server
   b. to keep a history of actions performed on the server
   c. to increase the administrator's already heavy workload
   d. to track software, updates, hardware, and drivers installed on the server
2. Why might it be important to record servers' interaction with one another?
   a. to see which other servers might be affected if this server stops functioning correctly
   b. to redirect all network traffic to the other servers
   c. so that you can move the RAID array to the most appropriate alternate
   d. to find an alternate for failover purposes
3. Why should you document the person performing the work?
   a. to properly assess blame for failed server administration
   b. to consult that person for more information about the work they performed on the server
   c. to verify that employees are being productive
   d. to make sure only that person works on the server in the future

4. When you see a server message that is difficult to interpret, you should do which of the following? (Choose two.)
   a. Ignore the message and tackle only the easier messages.
   b. Refer to preceding events.
   c. Access the vendor web site.
   d. Clear the error log to get rid of the message.
5. Error messages are:
   a. always critical—interpret every message
   b. seldom critical—only interpret if the server fails
   c. usually critical—some are innocuous as long as they do not affect performance or server availability
   d. worthy of attention only if other people notice server problems
6. When analyzing server problems, you should ask which of the following? (Choose all that apply.)
   a. Exactly what happened?
   b. Who does it affect?
   c. How does it happen?
   d. When does it happen?
7. Troubleshooting solutions are often as simple as:
   a. reprogramming the operating system
   b. an unplugged or loose cable
   c. a failed motherboard
   d. the human genome
8. One of the first things you should check when troubleshooting failed hardware is:
   a. the logs of the last few administrators who also worked on the server
   b. the BIOS version of the system
   c. the warranty
   d. the physical connection
9. Bend radius is:
   a. a reference to the degree to which a cable can be bent before losing signal integrity
   b. a yoga position
   c. a reference to the number of Token Ring hosts permitted in any given ring
   d. accidentally overwriting a recent driver version with an older one

10. What is likely to happen if you break the 5-4-3 rule?
    a. host name resolution problems
    b. excessive collisions and poor connectivity
    c. counting to infinity router loops
    d. SCSI ID conflicts
11. After replacing a server in the rack, it might be a good practice to:
    a. dust the case
    b. retract any anti-tipping mechanisms
    c. press in each connection
    d. run CHKDSK
12. An unusual noise in the power supply:
    a. could indicate imminent power supply failure
    b. can be ignored until you smell smoke
    c. indicates a fan problem only—you can continue using the power supply
    d. is not a problem provided the power supply is not also hot to the touch
13. If you suspect a problem with the NIC, how can you diagnose it?
    a. Use a packet sniffer to analyze each packet to see if the problem is consistent across all communications.
    b. Copy a large file to a network share, and see if a CRC check indicates that the file is corrupt.
    c. Use the NIC vendor's diagnosis utility.
    d. From a command prompt, type Ping 127.0.0.1.
14. Besides checking for proper voltages, how else can you use a multimeter?
    a. to check network cable amperages
    b. to check the speed of a cooling fan
    c. to check the resistance on thinnet cable
    d. to detect the location of a break in the cable
15. Which tool checks for breaks in the network cable?
    a. TDR
    b. PDQ
    c. DMZ
    d. PDC

16. Which of the following are valid loopback IP addresses? (Choose two.)
    a. 127.0.1.1
    b. 172.168.0.1
    c. 127.98.67.256
    d. 127.254.254.254
17. What can be sent to a WOL host to wake it up?
    a. a shock packet
    b. a magic packet
    c. a sleep packet
    d. any Ping packet
18. What do you use to create boot floppies for an OS/2 system?
    a. the MAKEBOOT utility
    b. the FORMAT A: /SYS command
    c. the INITIAL utility
    d. the Create Utility Diskette utility
19. The system periodically hangs without explanation. What might be a cause of this?
    a. failing power supply
    b. failing hard disk
    c. improper SCSI termination
    d. insufficient memory
20. One of the troubleshooting principles discussed in this chapter is:
    a. Troubleshoot as many problems at once as possible to save time.
    b. Troubleshoot only one issue at a time.
    c. Apply as many solutions at once as possible.
    d. When you finish troubleshooting, try to remember everything you did and write it down.

HANDS-ON PROJECTS

Project 12-1

In this project, you will observe individual drivers loading on a Windows NT 4.0, Windows 2000, or Windows .NET system.

1. Log on as Administrator.
2. Right-click **My Computer** and click **Explore**.

3. Double-click **C:\boot.ini**.

The Boot.ini file will not be visible unless you have deselected the option to hide protected operating system files. Select this option from Windows Explorer by selecting Tools, clicking Folder Options, and then deselect the option that appears on the View tab.

4. The Boot.ini file opens in Notepad.
5. Near the end of the file, you will see an entry that looks similar to this:[operating systems]multi(0)disk(0)rdisk(0)partition(1)\WINNT="Microsoft Windows 2000 Server" /fastdetect
6. After the /fastdetect item, add the **/SOS** switch.
7. Save the Boot.ini file and reboot the system.
8. Upon reboot, notice that you can see each driver as it loads. The drivers probably scroll by too fast to actually read each one, but if one of the drivers is problematic, the system will usually pause while trying to load it.
9. Allow the operating system to load normally.

Project 12-2

In this project, you will install Terminal Services on a Windows 2000 server.

1. Log on to a Windows 2000 Server computer as Administrator.
2. Add Terminal Services to the server: Click **Start**, point to **Settings**, and then click **Control Panel**.
3. From Control Panel, double-click the **Add/Remove Programs** item. The Add/Remove Programs dialog box opens. Double-click the net share, and then double-click the Win32 folder.
4. Click the **Add/Remove Windows Components** icon on the left of the interface.
5. Scroll down the list of Windows components. Click the **Terminal Services** checkbox and click **Next**.

Windows .NET servers are automatically configured for remote administration. You only add Terminal Server to run it in Application Server Mode.

6. Select to run the server in remote administration mode, and click **Next**.
7. Finish the wizard and, when prompted, reboot.

Be prepared to supply the Windows 2000 CD-ROM or enter a path to the Windows 2000 source files from which the operating system was installed.

8. Using Windows Explorer, locate C:\WINNT\system32\clients\tsclient\net and share the net folder with default permissions.
9. Close Windows Explorer.

Project 12-3

In this project, you will run a Terminal Services session to connect any Windows client to the Windows 2000 server you used in Project 12-2.

1. Log on to the domain from any Windows 9.x or later client.

If another computer is not available, you can perform all these actions on the Project 12-2 server. You will just connect the computer to itself in a Terminal Services session.

2. Browse Network Neighborhood or My Network Places (as it applies) and locate the server where the \net share exists.
3. Double-click the **net** share.
4. Locate and double-click the **Setup.exe** file. Click continue to proceed through the wizard.
5. Enter your name and organization, and click OK twice. Then, proceed to install the Terminal Services client with all default settings.
6. If so prompted, reboot the computer.
7. Log on again if necessary, and click **Start**, point to **Programs**, point to **Terminal Services Client**, and click **Terminal Services Client**. A dialog box like Figure 12-12 opens.

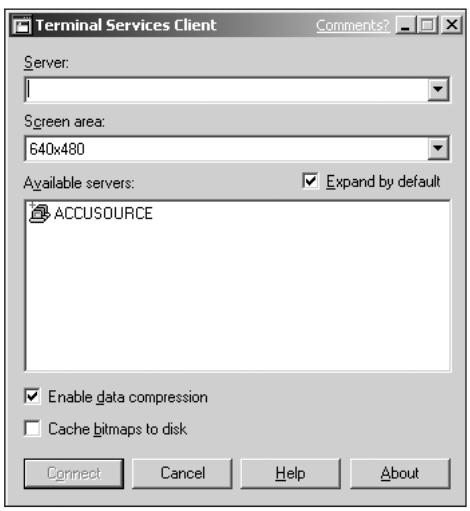

Figure 12-12 The Terminal Services Client dialog box

8. In the Server section, enter the IP address or host name of the Windows 2000 Terminal Services server.
9. Select a screen area resolution. Make sure that it is no larger than the actual resolution on the client where you are seated.
10. Click **Connect**. At the Windows 2000 logon screen, enter a user name and password and log on to the terminal server. Your desktop environment is virtually the same as if you were locally seated at the server. Browse around the interface—Start menu, Control Panel, and so on. The only difference you should see is perhaps a slightly diminished responsiveness to the mouse and keyboard (more obvious over a dial-up connection) and a limit of 256 colors. Open any window and leave it open.
11. Close the Terminal Services window. A message appears informing you that the session will continue to run even after you close the window and disconnect the session. Click OK to clear the message.
12. Access your former Terminal Services session by repeating Steps 7 through 10. Notice that the window you left open in Step 10 is still there—your session continued to run while you were "gone."
13. To actually close the session, click **Start**, **Shut Down**, and log off the session. It is now truly closed.

Project 12-4

In this project, you will view the Windows 2000 Event Log.

1. Log on as Administrator.
2. Right-click **My Computer** and click **Manage**. The Computer Management console opens.
3. In the left pane, expand the **Event Viewer** underneath System Tools. You should see three or more log categories, depending on the role of the server. The server in Figure 12-13 is a Domain Controller and DNS server, so it has a few more logs than might otherwise be present. Otherwise, your screen should look similar to this.
4. Select a log category in the left pane. Then, double-click any event in the right pane. Use the up and down arrows in the Event Properties dialog box to move up and down the list of messages.
5. Most servers probably have a warning or error of some kind. Make a note of a particular Event ID, search Microsoft's support site to see if you can find any information about the event, and answer the following questions:
   a. Does the event seem to be critical to the operation or performance of the server? How?
   b. Is there a resolution or additional information for the problem displayed in the Event Properties dialog box? If so, what is the resolution or information?
6. Close the Event Properties dialog box, and then close the Computer Management console.

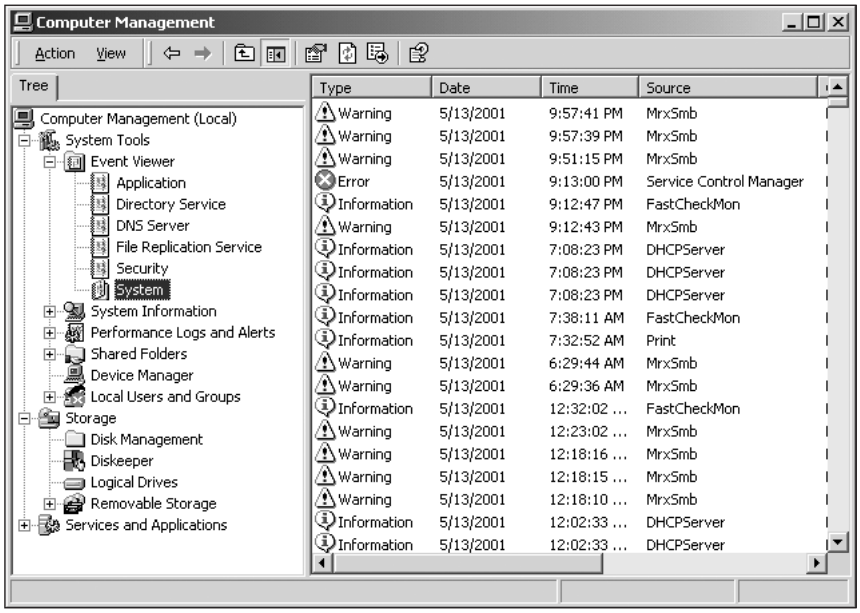

Figure 12-13 The Event Viewer categories with the System Event Viewer chosen

Project 12-5

In this project, you will search the Novell newsgroup to learn from others about Novell NetWare.

1. Using your web browser, access *www.novell.com*.
2. From the home page, click the **Get Support** link. Then select the Forums link from the support page.
3. On the Forums page, you can select the type of interface you want to use. If you have a newsgroup reader (such as Outlook Express), select **News Interface**. If you do not have a newsgroup reader or don't know, select **Web Interface**.
4. Select a link under support for **Operating System Products**, and then for **Server**, **NetWare 5.x**, and finally the **Utilities Forum**.
5. There are thousands of messages available, but only a few hundred have probably loaded into your newsgroup reader. Browse through messages of your choice and answer the following questions.
   a. What kind of troubleshooting problems had administrators been encountering?
   b. Were there responses that seemed to be incorrect? (This is a potential drawback to newsgroup support; sometimes people just offer a best guess.)
   c. What appears to be the best solution?
6. Close the browser and any open newsgroup readers.

Project 12-6

In this project, your instructor has deliberately caused a hardware problem on the server. Work with one or more other students to diagnose and troubleshoot the problem server. Be sure to utilize the troubleshooting logic from the chapter. Also, have someone record every step you take. Finally, when you have finished diagnosing and repairing the problem, write down on a separate paper what you think would be a useful server log record for this situation.

CASE PROJECTS

1. Kramer has been the supervising network administrator at a company for several years. Recently, Kramer noticed that several users were surreptitiously playing a game downloaded from somewhere on the Internet. Later, Kramer noticed an unusually high number of calls from users who said their files have started to disappear both locally and from the file server, and Kramer assigns an administrator to restore as many files as possible from tape backup. However, it disturbs him that so many files are missing. What would you say is a likely cause of the missing files? What should Kramer do? Would a written policy regarding programs that users can install on their computers help?
2. Jana, the IT supervisor for a medium-sized organization, returns from a two-week vacation to find that one of the company's two clustered web servers failed shortly after she left. The other administrators did not know how to fix the failed web server, and knowing that Jana would be able to fix it, they just let the problem remain unresolved. Jana looks at the server log records to see what happened to the server, and finds the last action was performed by Chuck the day after Jana left for vacation. Chuck's note only says: "Applied update. Started 2:14, ended 2:22." What else could Chuck have included in the log?